Networking Working Group
Internet-Draft
Intended status: Informational
Expires: January 3, 2013

JP. Vasseur, Ed.
J. Hui, Ed.
S. Dasgupta
G. Yoon
Cisco Systems, Inc
July 5, 2012

RPL deployment experience in large scale networks
draft-hui-vasseur-roll-rpl-deployment-01

Abstract

Low power and Lossy Networks (LLNs) exhibit characteristics unlike other more traditional IP links. LLNs are a class of network in which both routers and their interconnect are resource constrained. LLN routers are typically resource constrained in processing power, memory, and energy (i.e. battery power). LLN links are typically exhibit high loss rates, low data rates, are strongly affected by environmental conditions that change over time. LLNs may be composed of a few dozen to thousands of routers. A new protocol called the IPv6 Routing Protocol for Low-Power and Lossy Networks (RPL) has been specified for routing in LLNs supporting multipoint-to-point, point-to-multipoint traffic, and point-to-point traffic. Since RPL's publication as an RFC, several large scale networks have been successfully deployed. The aim of this document is to provide deployment experience on real-life deployed RPL-based networks.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

Status of this Memo

This Internet-Draft will expire on January 3, 2013.

Copyright Notice

Table of Contents

1. Introduction

Low power and Lossy Networks (LLNs) exhibit characteristics unlike other more traditional IP links. LLNs are a class of network in which both routers and their interconnect are resource constrained. LLN routers are typically resource constrained in processing power, memory, and energy (i.e. battery power). LLN links are typically exhibit high loss rates, low data rates, are strongly affected by environmental conditions that change over time. LLNs may be composed of a few dozen to thousands of routers.

A new protocol called the IPv6 Routing Protocol for Low-Power and Lossy Networks (RPL) has been specified for routing in LLNs supporting multipoint-to-point, point-to-multipoint traffic, and point-to-point traffic [RFC6550]. Since RPL's publication as an RFC, several large scale networks have been successfully deployed. The aim of this document is to provide deployment experience on real-life deployed RPL-based networks.

In addition to [RFC6550], companion documents have been defined that specify IPv6 packet options required for the proper operation of RPL, including [RFC6553] and [RFC6554].

This document makes use of the terminology defined in [I-D.ietf-roll-terminology].

RPL is a distance-vector routing protocol that builds a Destination Oriented Directed Acyclic Graph (DODAG) according to an Objective Function (OF). The OF is a defined set of rules that optimize paths against a set of metrics and/or constraints. A very basic OF, known as OF0, is specified in [RFC6552]. More involved OFs may be specified, such as the Minimum Rank with Hysteresis Objective Function specified in [I-D.ietf-roll-minrank-hysteresis-of].

Routing requirements documents spelled out in [RFC5673], [RFC5826], [RFC5548] and [RFC5867]) observe that it must be possible to take into account a variety of node metrics and/or constraints during path computation. Thus, a number of routing metrics and constraints for RPL have been specified in [RFC6551] for maximum flexibility according to the objectives and environment of the LLN.

RPL supports efficient loop detection using data-path route validation and supports both local and global route repair operations.

RPL makes use of the Trickle algorithm, which provides a density-aware mechanism for distributing and maintaining state within a network [RFC6206]. With simple local rules, the Trickle algorithm

adjusts the transmission period and suppresses redundant transmissions to minimize control traffic overhead in the steady state while propagating new information quickly. Trickle's suppression mechanisms ensures that control message overhead grows logarithmically with node density.

In maintaining point-to-multipoint routes, RPL supports two modes of operations: non-storing and storing. In both cases, the DODAG built by RPL according to the OF is used for hop-by-hop upstream routing towards the DAG Root. In non-storing mode, only the DAG Root maintains downward routes and all data packets must traverse the DAG Root. In storing mode, LLN routers also maintain downward routing state, allowing each LLN router to forward data packets to devices in their sub-DAG. LLN constraints, the network objective, and overall environment typically drives the choice of non-storing or storing mode and is left to the network administrator.

RPL, like many other routing protocols, is designed to be deployed in a number of different operational environments and [RFC6550] specifies a number of configuration parameters. Section 17 of [RFC6550] lists the following RPL constants and variables:

o DEFAULT_PATH_CONTROL_SIZE: This is the default value used to configure PCS in the DODAG Configuration option, which dictates the number of significant bits in the Path Control field of the Transit Information option. DEFAULT_PATH_CONTROL_SIZE has a value of 0. This configures the simplest case limiting the fan-out to 1 and limiting a node to send a DAO message to only one parent.

o DEFAULT_DIO_INTERVAL_MIN: This is the default value used to configure Imin for the DIO Trickle timer. DEFAULT_DIO_INTERVAL_MIN has a value of 3. This configuration results in Imin of 8 ms.

o DEFAULT_DIO_INTERVAL_DOUBLINGS: This is the default value used to configure Imax for the DIO Trickle timer. DEFAULT_DIO_INTERVAL_DOUBLINGS has a value of 20. This configuration results in a maximum interval of 2.3 hours.

o DEFAULT_DIO_REDUNDANCY_CONSTANT: This is the default value used to configure k for the DIO Trickle timer. DEFAULT_DIO_REDUNDANCY_CONSTANT has a value of 10. This configuration is a conservative value for Trickle suppression mechanism.

o DEFAULT_MIN_HOP_RANK_INCREASE: This is the default value of MinHopRankIncrease. DEFAULT_MIN_HOP_RANK_INCREASE has a value of 256. This configuration results in an 8-bit wide integer part of Rank.

o DEFAULT_DAO_DELAY: This is the default value for the DelayDAO Timer. DEFAULT_DAO_DELAY has a value of 1 second.

o DIO Timer: One instance per DODAG of which a node is a member. Expiry triggers DIO message transmission. A Trickle timer with variable interval in [0, DIOIntervalMin..2^DIOIntervalDoublings].

o DAG Version Increment Timer: Up to one instance per DODAG of which the node is acting as DODAG root. May not be supported in all implementations. Expiry triggers increment of DODAGVersionNumber, causing a new series of updated DIO message to be sent. Interval should be chosen appropriate to propagation time of DODAG and as appropriate to application requirements (e.g., response time versus overhead).

o DelayDAO Timer: Up to one timer per DAO parent (the subset of DODAG parents chosen to receive destination advertisements) per DODAG. Expiry triggers sending of DAO message to the DAO parent.

o RemoveTimer: Up to one timer per DAO entry per neighbor (i.e., those neighbors that have given DAO messages to this node as a DODAG parent). Expiry may trigger No-Path advertisements or immediately deallocate the DAO entry if there are no DAO parents.

Please refer to the .pdf version of this document to see the figures referred in further sections.

2. Objective of this document

Since its specification as a standard track RFC in March 2012, a number of RPL-based networks have been deployed in the field, some of small size, others of large scale. The aim of this document is to describe the successful deployment of a RPL-based LLN with 1,000 nodes. Other networks of even larger scale (5,000 to 10,000 nodes) are in progress and further revisions of this document will include their details.

It is nearly impossible to characterize the absolute performance of a protocol without looking at all the environmental factors and a large number of performance metrics. Furthermore such performance metric not only depends on the environment but also how the various protocol parameters have been configured. Similarly it would not make any sense to provide hard numbers on a performance characteristic of a protocol. For example, Open Shortest Path First (OSPF) routing protocol [RFC2328] may provide convergence times varying between few dozens of milliseconds to seconds depending on the network characteristics and protocol parameters. At one end of the spectrum, fast failure detection with fast Hellos or the use of other protocols such as Bidirectional Forwarding Detection (BFD) [RFC5880], combined with fast LSA generation, LSA prioritization, fast SPF triggering and an optimized SPF calculation (potentially combined with incremental SPF) would lead to a few dozens of milliseconds of convergence times. At the other end of the spectrum slow detection of failure, combined with low priority trigger of Link State Advertisement (LSA), poor implementation of the Shortest Path First (SPF) algorithm, long propagation delays, lack of LSA control plane packets may lead to convergence times of seconds!

While convergence time is not the critical performance metric in many LLN deployments, the convergence time example provided above is one that illustrates the challenge of providing performance results. This challenge generally applies to most other performance metrics.

As a result, the aim of this document is not to provide absolute performance numbers or parameter setting recommendations, but rather to share successful experience of the large scale deployment of RPL in a real-life deployment scenario.

To that end, we first provide several network characteristics such as the network topology, distribution of the link quality providing the link quality distribution according to the Expected Transmission count (ETX) link metric computed by the RPL nodes. Then we provide indications of how RPL was used in that particular network before showing several performance metrics observed in this network.

3. RPL Parameters Settings

This RPL network includes the following parameter settings:

o The Mode of Operation (MoP) is set to non-storing mode.

o Both local and global repair mechanisms are implemented. Note, however, because the network operates in non-storing mode, local repair simply poisons routes and does not create floating DAGs.

o MaxRankIncrease is set to 0, which significantly reduces the possibility of routing loops but also limits the capabilities of local repair.

o The OF is the Minimum Rank with Hysteresis Object Function using the ETX metric.

4. Network Characteristics

This network comprises of one thousand nodes and the distribution of the average hop counts is shown is Figure 1. This has been obtained by observing the topology (shown in Figure 2) for a period of 24 hours and tracking the hop count of all the nodes every 5 minutes. It can be seen that approximately 51% of the nodes are 1 hop away and 30% of nodes are 2 hops away on average. Another way of saying this is that approximately 81% of the nodes are 2 hops or less from the root. A snapshot of the network topology (the DODAG built by RPL) is depicted in Figure 2.

Figure 1: Distribution of average hop count of nodes observed over a 24-hour period.

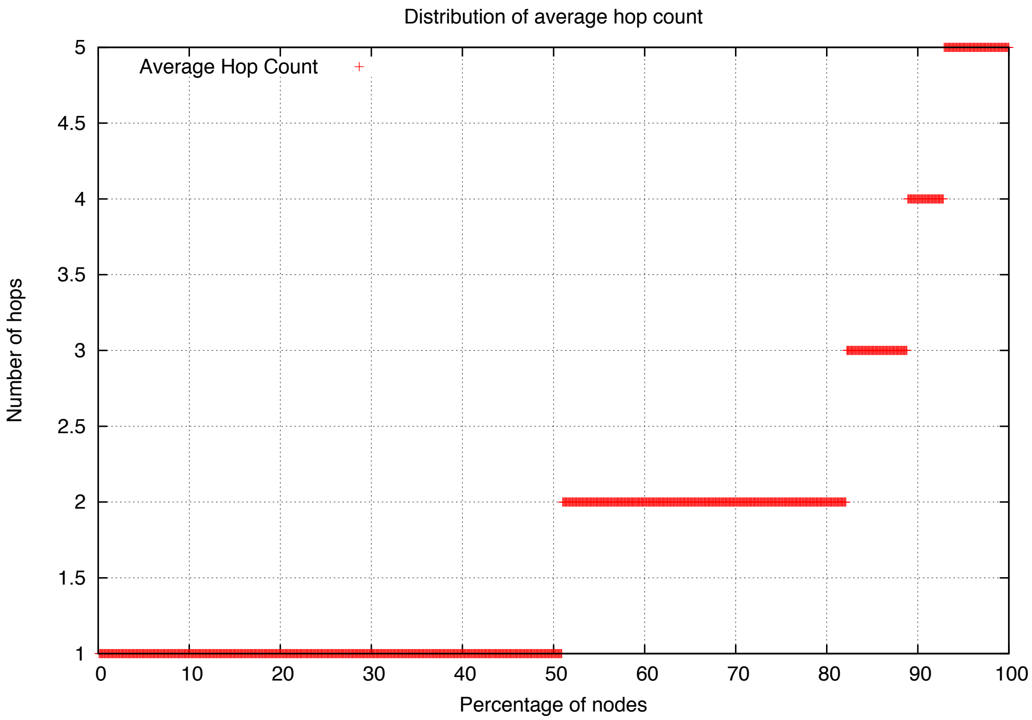

Figure 2: DODAG Topology built by RPL.

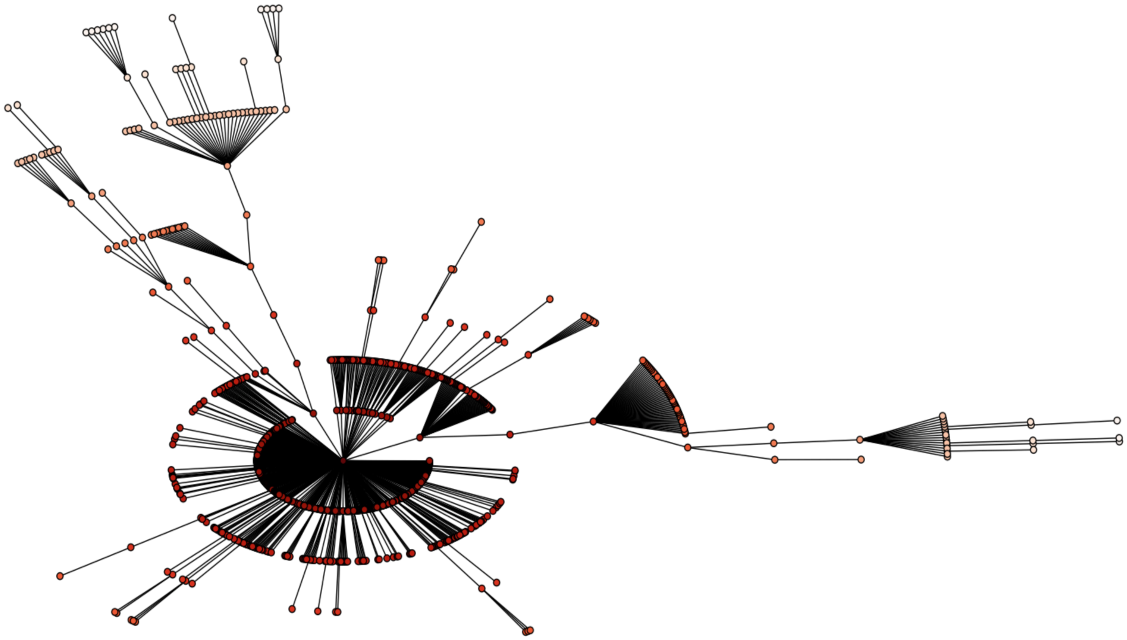

As with any LLN, one can observe that the some links are of good quality while others provides low path delivery rates: this can be seen by observing the link ETX in Figure 3. Note that we observed transient periods during which the ETX was much higher with links providing even intermittent connectivity (which is not always reflected in the ETX value due to the computation of the ETX is using moving averages to avoid network oscillations and over-reactions). Figure 3 was obtained by observing all the nodes periodically for a 24 hour period. We tracked the maximum and minimum ETX seen by the node as well. From the figure, we can see that almost 90% of the nodes had an average ETX of 1.25 or less over the 24 hour period.

Figure 3: Distribution of average, maximum and minimum ETX observed over a 24-hour period.

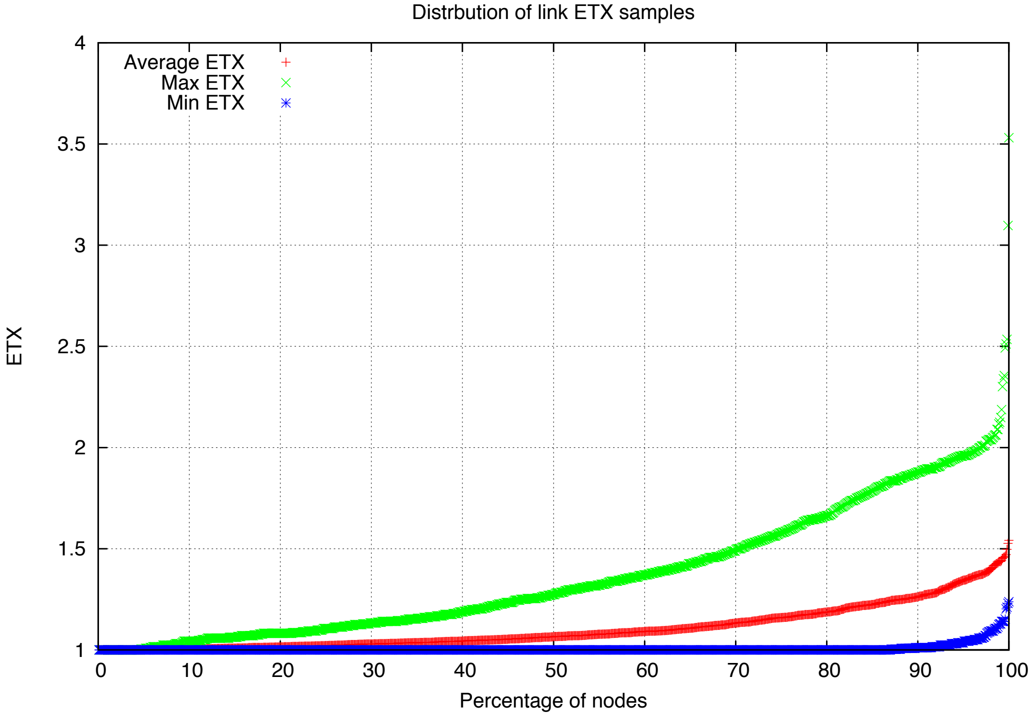

The LLN routers communicate using IEEE 802.15.4g links. Operating in the 902-928 MHz US ISM band, the links have an effective data rate of 75 kbps and employ frequency hopping to communicate across 64 channels with 400 kHz channel spacing.

In summary, it can be observed that the network is indeed a LLN, with lossy links.

The network is made of constrained nodes with limited processing power and available memory. The root slightly less constrained and main powered.

It is worth pointing out that the high density of this topology added stress on the routing protocol.

5. Performance Results

As pointed out in Section 1, there is not a single performance metric that could be provided to characterize the routing protocol performance.

Deep analysis of a number of network management events, logs on routers, and packet inspection operation have shown that the routing topology was quite stable even during unstable conditions. More importantly the observed packet delivery rate was always above 99%: by contrast with non LLN networks where the routing protocol is rarely responsible for non packet delivery because of the absence of routes to reach a destination, several LLN routing protocols have reported low delivery packet rates because of routing issues. In this particular network, the packet delivery rate was as high as 99% in all cases (link local packet retransmissions were handled by the IEEE 802.15.4 reliable link layer).

Since questions were raised in the past about the RPL control plane overhead although RPL has been designed for low overhead, we paid a particular attention to this performance metric. RPL has been designed to optimize the control plane overhead (thanks to the use of Neighbor Unreachability Detection (NUD) instead of routing hello packet to detect link/node failure, use of trickle algorithm for the transmission of the DIO packets, ...).

Thus we show in Figure 4 (to be published in a future revision) the RPL control traffic overhead (both the DIO and the DAO are shown in the network) relative to the available bandwidth provided by the links. Note that the RPL control plane traffic was observed on the most congested area of the network (on the DODAG root).

6. IANA Considerations

No action is required from IANA.

7. Security Considerations

This document provides informational data about existing deployments, thus security considerations do not apply.

8. Acknowledgements

The authors would like to acknowledge the contributions of Ibrahim Mortada for his very valuable contribution.

9. References

9.1. Normative references

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.

[RFC6550] Winter, T., Thubert, P., Brandt, A., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., Vasseur, JP., and R. Alexander, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks", RFC 6550, March 2012.

9.2. Informative references

[I-D.ietf-roll-minrank-hysteresis-of]
Gnawali, O. and P. Levis, "The Minimum Rank with Hysteresis Objective Function", draft-ietf-roll-minrank-hysteresis-of-11 (work in progress), June 2012.

[I-D.ietf-roll-terminology]
Vasseur, J., "Terminology in Low power And Lossy Networks", draft-ietf-roll-terminology-06 (work in progress), September 2011.

[RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.

[RFC5548] Dohler, M., Watteyne, T., Winter, T., and D. Barthel, "Routing Requirements for Urban Low-Power and Lossy Networks", RFC 5548, May 2009.

[RFC5673] Pister, K., Thubert, P., Dwars, S., and T. Phinney, "Industrial Routing Requirements in Low-Power and Lossy Networks", RFC 5673, October 2009.

[RFC5826] Brandt, A., Buron, J., and G. Porcu, "Home Automation Routing Requirements in Low-Power and Lossy Networks", RFC 5826, April 2010.

[RFC5867] Martocci, J., De Mil, P., Riou, N., and W. Vermeylen, "Building Automation Routing Requirements in Low-Power and Lossy Networks", RFC 5867, June 2010.

[RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", RFC 5880, June 2010.

[RFC6206] Levis, P., Clausen, T., Hui, J., Gnawali, O., and J. Ko, "The Trickle Algorithm", RFC 6206, March 2011.

[RFC6551] Vasseur, JP., Kim, M., Pister, K., Dejean, N., and D. Barthel, "Routing Metrics Used for Path Calculation in

Low-Power and Lossy Networks", RFC 6551, March 2012.

[RFC6552] Thubert, P., "Objective Function Zero for the Routing Protocol for Low-Power and Lossy Networks (RPL)", RFC 6552, March 2012.

[RFC6553] Hui, J. and JP. Vasseur, "The Routing Protocol for Low-Power and Lossy Networks (RPL) Option for Carrying RPL Information in Data-Plane Datagrams", RFC 6553, March 2012.

[RFC6554] Hui, J., Vasseur, JP., Culler, D., and V. Manral, "An IPv6 Routing Header for Source Routes with the Routing Protocol for Low-Power and Lossy Networks (RPL)", RFC 6554, March 2012.

Authors' Addresses

JP Vasseur (editor)
Cisco Systems, Inc
11, Rue Camille Desmoulins
Issy Les Moulineaux, 92782
France

Email: jpv@cisco.com

Jonathan Hui (editor)
Cisco Systems, Inc
560 McCarthy Blvd.
MILPITAS, CALIFORNIA 95035
UNITED STATES

Email: johui@cisco.com

Sukrit Dasgupta
Cisco Systems, Inc
300 Beaver Brook Road
BOXBOROUGH, MASSACHUSETTS 01719
UNITED STATES

Email: sukdasgu@cisco.com

Giyoung Yoon
Cisco Systems, Inc
560 McCarthy Blvd.
MILPITAS, CALIFORNIA 95035
UNITED STATES

Email: giyoon@cisco.com